

Mitigating Redundant Data Transfers for Mobile Web Applications via App-Specific Cache Space

Yun Ma, Xuanzhe Liu, Shuailiang Dong, Yunxin Liu, Tao Xie

Abstract

Redundant transfer of resources is a critical issue for compromising the performance of mobile Web applications (a.k.a., apps) in terms of data traffic, load time, and even energy consumption. Evidence shows that the current cache mechanisms are far from satisfactory. With lessons learned from how native apps manage their resources, in this paper, we propose the ReWAP approach to fundamentally reducing redundant transfers by restructuring the resource loading of mobile Web apps. ReWAP is based on an efficient mechanism of resource packaging where stable resources are encapsulated and maintained into a package, and such a package shall be loaded always from the local storage and updated by explicitly refreshing. By retrieving and analyzing the update of resources, ReWAP maintains resource packages that can accurately identify which resources can be loaded from the local storage for a considerably long period. ReWAP also provides a wrapper for mobile Web apps to enable loading and updating resource packages in the local storage as well as loading resources from resource packages. ReWAP can be easily and seamlessly deployed into existing mobile Web architectures with minimal modifications, and is transparent to end-users. We evaluate ReWAP based on continuous 15-day access traces of 50 mobile Web apps that suffer heavily from the problem of redundant transfers. Compared to the original mobile Web apps with cache enabled, ReWAP can significantly reduce the data traffic, with the median saving up to 51%. In addition, ReWAP can incur only very minor runtime overhead of the client-side browsers.

1 Introduction

Redundant transfer of resources¹ refers to the case where a previously fetched resource is downloaded again from the network before the resource is actually updated. For mobile Web applications (a.k.a. apps) [1], redundant transfers remain as a critical performance issue leading to duplicated data transmission, long page load time, and high energy drain [2][3].

Redundant transfers originate from apps' resource-management mechanism that is to determine whether a resource should be loaded locally or remotely. Web cache is a conventional resource-management mechanism adopted by Web apps. Web developers can configure cache policies, such as expiration time and validation flag, on those resources that are likely to be loaded from the local storage. The browser maintains a cache space and deals with the cache logic for all the Web apps running in it. However, our previous work [4] found that there is a big gap between the ideal and actual cache performance of mobile Web apps. For example, for the mobile versions of top-100 websites of Alexa, although more than 70% of resources can be loaded from the cache when these websites are revisited after one day, less than 50% of these cacheable resources are actually loaded from the cache. Surprisingly, all resource transfers are redundant for some well-known websites when they are revisited after one day. We also revealed two major causes for redundant transfers: (1) the imperfect cache configuration, such as heuristic expiration and conservative expiration time; and (2) the undesirable Web development practice, such as using random strings to name resources for enforcing their refresh.

Due to the dynamics of mobile Web apps, it is difficult for Web developers to properly configure the apps' cache policies. Short expiration time may lead to redundant transfers, while long expiration time

¹In this paper, resources refer to resource objects constituting an app (such as HTML, JavaScript, CSS, and images of a Web app; native code, media files, and layout files of a native app)

may result in the usage of stale resources. As a result, using cache policies is not a desirable mechanism to accurately determine whether resources should be loaded locally or remotely [5].

To fundamentally reduce redundant transfers for mobile Web apps, the resource management in native apps can provide some useful inspirations. Resources of native apps are managed directly and explicitly by app-specific logics to control where to load resources and when to update the local resources. Intuitively, native apps explicitly distinguish their static resources from dynamic ones, and encapsulate the static resources into a package that is installed into a dedicated space allocated by the underlying operating system. When a native app is running, the app logic controls that only its dynamic resources are downloaded on demand to provide the fresh data to users, while the static resources in the installed package are always fetched locally. When the static resources have to be updated, a new resource package is downloaded and installed to refresh all the static resources.

However, there are two main challenges for Web developers to adopt such a package-based resource management specific to a Web app. First, it is hard to maintain the resource package. Resources of Web apps are loosely coupled, and usually updated independently and casually without influencing each other. As a result, it is tedious and error-prone to decide which resources should be put into the package and when to update the package. Second, it is hard to enable mobile Web apps to use the package. Modern Web apps are complex and there are many mature Web development frameworks. As a result, a lot of manual efforts have to be needed to realize or refactor mobile Web apps to benefit from the package-based resource management.

To address these challenges, in this paper, we propose the *ReWAP* approach to restructuring mobile Web apps to be equipped with package-based resource management while requiring minimal developer efforts. The key rationale of ReWAP is to provide more efficient and app-specific control of resource management rather than relying on only the current mechanisms such as Web cache, to avoid the caused unnecessary redundant resource transfers. By retrieving the update of resources of mobile Web apps, ReWAP automatically maintains resource packages by accurately identifying which resources should be loaded from the local storage for a considerably long period. Based on the package information, ReWAP automatically checks the update of resource packages, refreshes the resources in the package when the resource package is updated, and loads resources from resource package for mobile Web apps. To integrate ReWAP with existing mobile Web apps, Web developers need only minor modifications to their existing apps. In summary, *ReWAP shares the same spirit of resource management mechanisms as those in the installation package of native apps but in the way of Web.*

To the best of our knowledge, our work is the first to facilitate Web developers to effectively reduce redundant transfers of mobile Web apps by conducting resource management in a similar way as native apps. More specifically, this paper makes the following main contributions:

- We design ReWAP, a packaging approach for a mobile Web app to accurately identify resources that can be loaded from the local storage for a considerably long time. The maintained package can maximize the benefit of data-traffic saving by considering all the users of the mobile Web app.
- We implement ReWAP to minimize developer efforts of restructuring existing mobile Web apps. Web developers can easily integrate ReWAP in their current implementation of a mobile Web app, and the end-users are completely unaware of the existence of ReWAP when they access the Web apps.
- We conduct experiments based on 15-day access logs of 50 mobile Web apps that suffer heavily from redundant transfers to evaluate the effectiveness of ReWAP. Compared to the original mobile Web apps with cache enabled, ReWAP can significantly save the data traffic with the median value up to 51% and the maximum of almost 100%. In addition, ReWAP incurs only quite small runtime overhead of the client-side browsers.

The remainder of this paper is organized as follows. Section 2 illustrates the problem of redundant transfers with an example and compares the resource-management mechanisms of Web apps and native apps. Section 3 presents the overview of the ReWAP approach. Sections 4 and 5 show the details

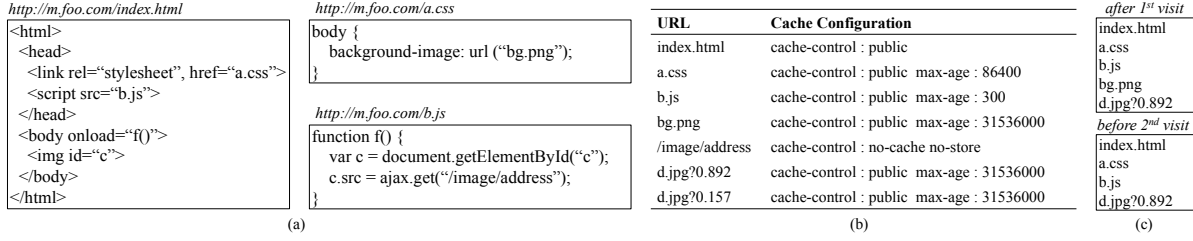


Figure 1: Motivating example. (a) Resource excerpts of a Web app; (b) Cache configuration; (c) Cache entries between two visits.

of ReWAP’s key components, i.e., Package Engine and Wrapper, respectively. Section 6 describes the implementation of ReWAP and demonstrates its easy deployment. Section 7 evaluates ReWAP based on top Web apps of Alexa. Section 8 discusses the related work and Section 9 concludes the paper.

2 Background and Motivation

In this section, we present the background and motivation for leveraging the resource-management mechanisms of native apps to improve Web apps. We first describe a motivating example to illustrate the problem of redundant transfers. Then we compare the resource-management mechanisms of Web and native apps.

2.1 Redundant Transfer in Mobile Web Apps

Although the resource management of Web apps is flexible enough to achieve easy-to-access and always-updated features, it could lead to redundant transfers of resources. We illustrate redundant transfers via an example shown in Figure 1.

Figure 1(a) are resource excerpts of a mobile Web app “http://m.foo.com/”. The HTML resource indicates that the app includes a layout resource “a.css” and a JavaScript resource “b.js”. When “a.css” is being evaluated, a background image “bg.png” is identified. After parsing the HTML resource is finished and the `onload` event is triggered, JavaScript function `f` is executed to get an address of an image by requesting a service “/image/address”. Suppose that the returned address is “d.jpg?0.892”, and then the image is retrieved. Therefore, when the app is visited at the first time, 6 resources are actually retrieved. Figure 1(b) shows the cache configuration of these resources. The HTML is not configured with an explicit expiration time so the browser assigns a random time that is usually not very long, e.g., 30 minutes. The expiration time of CSS, JavaScript, and images are configured as 1 day, 5 minutes, and 1 year, respectively. The service “/image/address” is configured as no-cache and no-store to ensure obtaining the latest address at every visit. The top table in Figure 1(c) shows the resources in the browser cache after the first visit.

Assume that the app is revisited after one hour and all the related resources have not been updated except the “/image/address”. The bottom table in Figure 1(c) shows the cached resources before the second visit. It can be seen that the background image “bg.png” has been removed out of the cache due to the limited size of cache on mobile devices because all the Web apps accessed by a browser share a fixed size of cache space.

Given the current status of cache, at the revisit of this app, several resources that could have been loaded from the cache are actually re-downloaded from the network, leading to redundant transfers of resources falling in the following main categories [2, 4].

RT1: Resources that are moved out of the cache.




Figure 2: Resource management of (a) Web apps and (b) native apps.

Due to the imperfect implementation of cache on mobile browsers such as limited size and non-persistent storage, resources in the cache may be removed out of the cache after some time. In the preceding example, the background image “bg.png” is removed out of the cache and has to be re-downloaded when the Web app is revisited.

RT2: Resources that are wrongly judged as expired.

Each resource has to be configured by developers with a cache policy. Due to the imperfect cache configuration of resources whose expiration time is either configured to be too short or not configured but assigned heuristically by browsers, many resources are falsely judged by browsers as expired ones, and have to be validated or re-downloaded. In the example, the HTML resource has not been assigned an explicit expiration time, and the expiration time of the JavaScript resource is configured to be too short. As a result, these two resources cannot be loaded from the local environment when the Web app is revisited.

RT3: Resources that are requested by new URLs but have the same content with cached ones.

Resource Loader of browsers uses URLs to uniquely distinguish among resources. Resources with different URLs are regarded as totally different ones. In the example, the same image “d.jpg” has different URLs at two visits, resulting in being fetched twice. URL changing is usually adopted to realize backend load balance according to URL routing. Although such a practice can improve the performance of backend servers, it actually harms the loading process of mobile Web apps.

2.2 Resource Management of Web Apps and Native Apps

A key reason for the preceding redundant transfers is the inefficient resource management of Web apps. Figure 2(a) illustrates the resource-management mechanism of Web apps. Web apps rely on the underlying browsers to manage their resources. All Web apps in a browser share a common cache space whose size is usually small on mobile devices. When a user launches a Web app in the browser (①), resources for rendering the Web app are dynamically identified and all the resource requests are handled by the Resource Loader component in the browser (②). Based on the Web cache mechanism [6], the Resource Loader determines whether to load the resource from the cache (③) or downloads it from the server (④). After retrieving the resource, the Resource Loader returns it to the Web app (⑤). In summary, Web apps rely on the app-independent browser logics to manage resources. Such a mechanism makes Web apps flexible for resource management so that Web apps can be always up-to-date. However, Web apps cannot have the full control of resources to be loaded from the local storage and when to update the local resources. As a result, redundant transfers arise when the cache policies are not configured properly or the browser removes cached resources.

In contrast, the resource management of native apps works in a different fashion and can be more efficient. Figure 2(b) illustrates the resource-management mechanism of native apps. Native apps separate resources into two sets, i.e., the dynamic resource set and the static resource set. Static resources are encapsulated into a resource package. Before using the native app, the resource package has to be installed on the device. When a user launches the app (①), the App Logic controls to load static resources from the App Space (②) and dynamic resources from the server (③). Usually, there is a built-in Update Manager for updating the static resources. The Update Manager checks update with the server (④) in some situations (e.g., every time when the app is launched) to find whether the resource package has been updated (⑤). If a new package is retrieved, the Update Manager confirms with the users whether to update the app (⑥). If agreed (⑦), then the Update Manager refreshes the static resources with the new resource package (⑧). In summary, native apps have app-specific logics to control the resources loaded from the local environment and the update of local resources.

Comparing the two resource-management mechanisms reveals that native apps can manage their resources based on app-specific logic with resource packages while Web apps cannot precisely manage their resources. The insight underlying our new approach is that redundant transfers originate from the principle of the resource-management mechanism adopted by Web apps.

3 Approach Overview

To fundamentally reduce the redundant transfers, we propose our new solution with the key rationale of lessons learned from the resource-management mechanism used by native apps, while keeping the advantages of the mechanism used by Web apps. More specifically, mobile Web apps can encapsulate stable resources into a package and make the resources in the package always loaded locally rather than being fetched from the servers, while other resources are regularly loaded by the browser’s default mechanism. All the resources in the package are refreshed together also by the default mechanism only when the resource package get updated. The update of package should also follow the way of the Web without the intervention of end users.

To this end, we propose the ReWAP approach to restructuring mobile Web apps to be equipped with package-based resource management. ReWAP can accurately identify the resources that should be loaded from the local storage for a considerably long time and that can be refreshed together with minimal cost when the package is updated. Other than the native apps, such a packaging mechanism follows the conventional way of the Web, i.e., the updating and refreshing of packaged resources still use the browser’s default cache mechanism. The Web developers can simply integrate ReWAP into their existing mobile Web apps with only minor modifications. As we will show later, they need only to redirect the entrance of the app to a Wrapper that delegates the resource loading. Meanwhile, the client-side browser performs regularly without additional modifications.

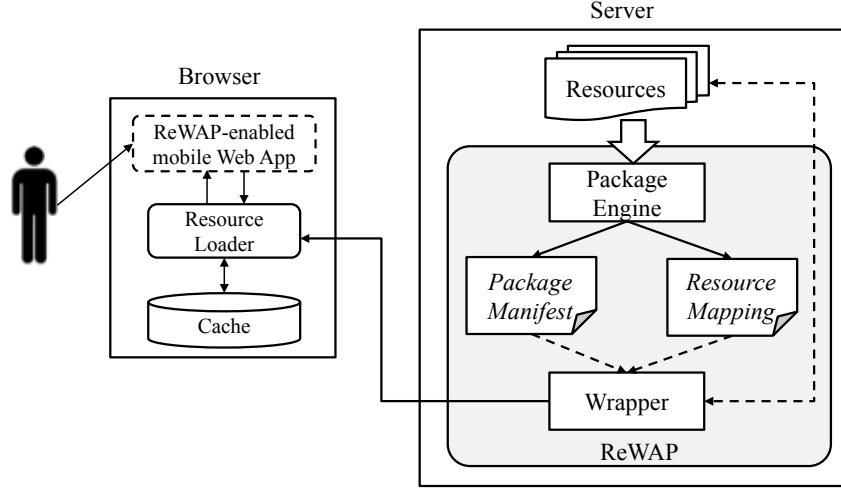


Figure 3: The ReWAP approach.

Figure 3 shows the overview of the ReWAP approach. In general, ReWAP can be deployed at the server of Web apps. It consists of a Package Engine to automatically generate and maintain resource packages, and a Wrapper for mobile Web apps to use and update resource packages in the local storage. Each Web app has a unique resource package for all their users, just like the installation package of a native app. Note that a mobile Web app can consist of many Web pages. Our approach currently treats each Web page independently. As different Web pages of a Web app may share similar resources [7], it is possible to further improve the performance by packaging all the Web pages of a Web app together. Such optimization is to be figured out as future work. In the rest of this paper, the term “Web app” denotes a single Web page referenced by an HTML document.

By retrieving the update of resources constituting a mobile Web app, the Package Engine generates and maintains a resource package with two configuration files: Package Manifest and Resource Mapping. The Package Manifest specifies which resources are in the resource package. The update of Package Manifest indicates the update of the corresponding resource package. The Resource Mapping keeps the relationship between URLs and unique resource entities. Resources that have the same content but are identified by different URLs are mapped into one resource entity according to Resource Mapping. Therefore, the generated package is highly accurate to cover more resources.

The Wrapper is essentially a separate HTML page where Web developers can easily enable their mobile Web apps with the package-based resource management. When a ReWAP-enabled mobile Web app is launched, the Wrapper is first fetched from the server. Then the Wrapper controls the loading process on the browser (we use dotted lines to represent the flow taking place on the client side). The Wrapper checks whether the resource package has been updated according to the Package Manifest. If updated, all resources in the package are refreshed and stored into an App-Specific Space according to Resource Mapping. After the refreshing, the Wrapper loads the app. The Wrapper intercepts all the resource requests to determine whether to load a resource from the App-Specific Space or as usual based on the Package Manifest.

ReWAP is deployed as a service on the same server with the target mobile Web app. For example, to integrate ReWAP with the motivating mobile Web app in Section 2, the developer can specify the app’s URL <http://m.foo.com/index.html> and then launches ReWAP service on the server m.foo.com. The Package Engine is then automatically started as a background process at the server side, while the Wrapper is also generated on the server with a URL, e.g., <http://m.foo.com/index/wrapper.html>. At last, the developer configures the server, making the requests to [index.html](http://m.foo.com/index.html) redirected to the URL of




Figure 4: The workflow of the Package Engine.

the Wrapper. When a user visits the ReWAP-enabled app, the Wrapper is loaded first to the browser, dealing with resource packages. Then the Wrapper loads the index.html by an AJAX call and intercepts all the resource requests. On the whole, deploying ReWAP requires only minimal modifications to existing mobile Web apps.

In the next two sections, we present the technical details on how the Package Engine maintains the resource package and how the Wrapper supports resource package at the runtime of Web apps.

4 The Package Engine

Figure 4 shows the four phases to maintain the resource package. In the first phase *Retrieve Resources*, the Package Engine repeatedly visits the Web app via a real browser runtime (e.g., WebKit) that is deployed at server-side, retrieves all the needed resources to render the app, and stores them into the Resource Repository. The visit is performed with a given time interval or every time when informed by the server. In the second phase *Normalize Resources*, based on the resources in the Resource Repository, resources that have the same content but identified with different URLs are normalized as a unique resource that is not required to be re-downloaded. The normalized resource is identified based on a regular expression. Resources whose URLs can match the regular expression are treated as the same resource. A Resource Mapping file is generated to maintain the relationship between normalized and original resources. In the third phase *Predict Update Time*, according to the history of resource updates kept in the Normalized Resource Repository, the Package Engine predicts the update time of each normalized resource. In the last phase *Generate Package*, a subset of normalized resources is selected to generate the Package Manifest based on the benefit brought to the mobile Web app. Since all the users of an app share the same package, we use the averagely saved data traffic as the main metric to quantify the benefit. The generation takes place in two conditions. If the current resource package is invalid, i.e., any resource in the package has been changed, then a passive selection is performed to select the resources with the highest benefit to generate the new resource package. Otherwise, if the current package is still valid, an active selection can generate a new package only when the benefit of the new package exceeds the old one over a given

threshold. We next present the details of each phase, respectively.

4.1 Retrieving Resources

In the first phase, by periodically revisiting the target Web app, we can retrieve the update of the Web app’s resources. We define a “*concrete resource*” structure to represent each retrieved resource. A “concrete resource” has the following fields: (1) *URL*, which is the identifier of the resource; (2) *MD5*, which is the checksum of the resource content; (3) *size*, which is the length of the resource content; (4) *cache duration*, which is the configured expiration time. All these fields can be obtained from the HTTP response message. We define R_t as the set of concrete resources retrieved at time t .

In order to ensure the completeness of the retrieved resources, we employ a real browser facility to actually launch and render the Web app. Otherwise, a lot of resources may be missed by only parsing the HTML document. We record all the HTTP traffic occurred by the browser facility in the progress of visiting a Web app.

Here, an important issue is the timing to retrieve resources because we have to ensure that the resource status is always up-to-date to avoid inconsistencies. One possible solution is to make the server notify the Package Engine at the time when the resources are updated so that resource updates can be exactly captured. Another solution is to trigger by a retrieving frequency. The retrieving frequency can be self-adjusted according to the features of the target Web apps. For example, apps that change more often can be revisited more frequently. In our current design, we use a fixed retrieving frequency to trigger resource retrieving. We leave the server-informed mechanism and self-adjusted frequency to the future work.

4.2 Normalizing Resources

In the second phase, the Package Engine identifies the resources that have different URLs but the same content at different visits to the Web app. We denote this kind of resources as “*cross-dress resources*”. We normalize cross-dress resources into one normalized resource. We keep the relationship of the normalized resource and one concrete resource in the Resource Mapping. The final Package Manifest consists of normalized resources so that cross-dress resources do not have to be re-downloaded for multiple times.

We observe that there are two frequent URL patterns of the cross-dress resources. One is the query strings generated by JavaScript, e.g. *Math.random()*, or by server scripts. In the motivating example, the URL of the image “d.jpg” has two different query strings “?892” and “?157” at two visits but the image does not change. In such a case, the URLs vary only in the query, i.e., the random value. The other pattern is the *CDN prefixes*. At different visits, resources could be retargeted at different CDN servers. In such a case, the paths of the URL are the same but the domain part of the URL could be changed according to the target CDN servers. Based on the two patterns, we assume that the URLs of cross-dress resources can be different in a certain part of the URL. Therefore, we can apply the *Longest-Substring* algorithm to find the base string of the URLs and use a regular expression to represent the changing part. For example, the image “d.jpg” in the motivating Web app has two different query strings “?892” and “?157” at two visits. So we can use the regular expression “*d.jpg\{?*\}*” to represent the normalized image.

We define a “*normalized resource*” structure to represent a unique normalized resource. A normalized resource is generated by aggregating cross-dress resources. It has all the fields of the “*concrete resource*” structure. The additional fields of “*normalized resource*” include: (1) *expression*, which describes the URL pattern of cross-dress resources; (2) *predicted time*, which describes the estimated duration time that the resource remains unchanged; and (3) *status*, which records all the historic statuses of the resource. We use $status_t$ to denote the status at time t . Each status can be “*inexistent*”, “*changed*”, or “*unchanged*”. Such historic status information is used to predict the update time. We define H_t as the set of normalized resources at time t . H_t is updated every time when concrete resources are retrieved.

We should enforce a one-to-one mapping between the normalized resources and the concrete resources of each visit in order to prevent resources with different contents from being matched to one normalized

resource. To this end, we compare the MD5 of the resource content to determine whether two resources are the same.

<p>Input: Last set of normalized resources H_{t-1}, current set of concrete resources R_t</p> <p>Output: Updated set of normalized resources H_t</p> <pre> 1 INITIAL $H_t \leftarrow H_{t-1}$; 2 foreach $h \in H_t$ do 3 INITIAL $h.status_t \leftarrow "inexistent"$; 4 end 5 foreach $r \in R_t$ do 6 $P \leftarrow FindSameURL(H_t, r)$; 7 $q \leftarrow FindSameMD5(H_t, r)$; 8 if $q \neq null$ then 9 $q.expression \leftarrow CalRegExpr(q.expression, r.URL)$; 10 $q.status_t \leftarrow "unchanged"$; 11 end 12 else if $P.size = 1$ then 13 $P.status_t \leftarrow "changed"$; 14 $UpdateResource(P)$; 15 end 16 else 17 $RemoveResource(P)$; 18 $AddResource(r)$; 19 end 20 end 21 $CheckMapping(R_t, H_t)$; 22 return H_t. </pre>

Algorithm 4.1: Normalize resources.

Algorithm 4.1 describes the process of managing the set of normalized resources. Given the last set of normalized resources H_{t-1} , and the current set of concrete resources R_t , the algorithm returns the updated set of normalized resources H_t . We assign H_{t-1} to H_t at first and initialize all resources' status of time t as "inexistent" (Lines 1-4). Then we add, update, and remove resources in different cases (Lines 5-20). Additionally, we need to handle conflicts in the new set of normalized resources to ensure the one-to-one mapping (Line 21).

4.3 Predicting Update Time of Resources

In the third phase, we infer whether a resource is sufficiently stable by predicting the update time of the resource. We assume that the evolution history of a resource can reflect the trend of resource updates. For example, if a resource is updated every day in the history, it is likely to be updated in the next day. Therefore, we design an algorithm to predict the update time of resources based on their evolution histories.

Algorithm 4.2 shows the details of how our prediction works. By examining the evolution histories of some resources, we find that after a resource disappears at one visit, the possibility of its reappearance is rather small. So every time when an "inexistent" status is captured for a resource, we immediately set its predicted time to 0 (Lines 1-3). For other resources, we mainly capture the total times of "changed" status and predict the next time when the resource is likely to change. In some cases, the resource can update in an unusual fashion, so we should not aggressively change the predicted time. For example, if a resource is updated once every day in the history and at one time it is updated one hour after the last update, we should use a modest way to reduce the predict time. Here, we use the gradient descent algorithm [8]

to predict the resource update time (Line 5). Furthermore, if no “changed” status is captured, we set the predicted time according to the numbers of “unchanged” other than infinite (Lines 6-8). Finally, we remove all the resources whose predicted time is 0 in order to limit the numbers of historic resources (Lines 10-12).

Input: Historic status $status_0, \dots, status_t$ of a normalized resource $h \in H_t$, visiting interval vi
Output: Predicted update time of h

```

1 if  $h.status_t = \text{"inexistence"}$  then
2   |  $h.predictedtime \leftarrow 0$ ;
3 end
4 else
5   |  $h.predictedtime \leftarrow GDM(status_0, \dots, status_t)$ ;
6   | if  $h.predictedtime = inf$  then
7     |  $h.predictedtime \leftarrow |status.unchanged| * vi$ ;
8   | end
9 end
10 if  $h.predictedtime = 0$  then
11   |  $RemoveResource(h)$ ;
12 end
```

Algorithm 4.2: Predict update time of normalized resources.

4.4 Generating Package

After estimating the update time of each resource, we can select the stable resources that can be packaged. We regard that resources in the package can bring some benefit to the Web app. We assume that all the Web app’s users share a common resource package. Given that different users may revisit the Web app at different times and with different frequencies, the benefit varies among different users. To measure the overall benefit that can be gained by the packaged resource, we assume a user distribution function σ to represent the percentage of users at different revisiting intervals. We define a metric, the average saved data traffic, to quantify how much data traffic all the Web app’s users can save on average given the user distribution function.

Let us assume that a subset of resources $M \subset H_t$ is selected. Suppose that T is the minimum predicted time in M , we then have an expectation that such a resource package M will be updated at time T . For each normalized resource in M , the browser does not need to request this resource before T . On the contrary, without the resource package, each resource should be requested from the server after its cache duration. Thus, for each resource, if T exceeds the cache duration, the traffic saving comes from the difference between the configured cache duration and our predicted time. If the predicted time is less than the cache duration, the resource can still be loaded from the cache according to our package mechanism, incurring no extra traffic.

Algorithm 4.3 shows how to select the best resource package. We first sort the normalized resources according to the predicted time in ascending order (Line 1). Among all subsets of resources whose minimum predicted time is T , the benefit of a smaller set cannot exceed that of a bigger one. Thus we do not need to enumerate all potential packages. We enumerate the potential T (Lines 2-10), and calculate the benefit of the largest set whose minimum predicted time is T (Lines 5-10). Finally, we choose the package that provides maximum benefit (Lines 11-13).

To make the resource package sufficiently stable, we may not always use the resource package with the largest benefit. We first check whether the latest resource package is still valid where all the resources in the latest package has not been updated. If any resource is changed, we just use the package generated by Algorithm 4.3 to replace the invalid package. If the latest resource package is still valid, then we replace the latest package only when the benefit of the new package exceeds the current one by a given

```

Input: Current set of normalized resources  $H_t$ , user distribution  $\sigma$ 
Output: Resource package  $M$ 
1 Sort  $H_t$  based on its predicted time in ascending order;
2 for  $i \leftarrow 0$  to  $|H_t|$  do
3    $benefit(i) \leftarrow 0$ ;
4    $T \leftarrow H_i.predictedtime$ ;
5   for  $j \leftarrow i$  to  $|H_t|$  do
6     if  $H_j.cacheduration < T$  then
7        $benefit(i) + = \sigma(H_j.cacheduration, T) * H_j.size$ ;
8     end
9   end
10 end
11 Select  $i$  where  $benefit(i)$  is the largest;
12  $M \leftarrow H_t(i, i + 1, \dots, |H_t|)$ ;
13 return  $M$ .

```

Algorithm 4.3: Select the packaged resources.

threshold.

5 The Wrapper

The functionality of the Wrapper is to equip mobile Web apps with the ability to use the resource packages atop the resource-management mechanism of Web browsers. Figure 5 shows the workflow of the Wrapper. The Wrapper has an App-Specific Space, which is a dedicated local storage to store packaged resources for each Web app. Each Web app has its own App-Specific Space that is not shared with other Web apps.

When an end user visits the ReWAP-enabled mobile Web app, the Wrapper is first loaded from the server and runs in the browser. It has three phases to load the target mobile Web app.

Checking Package Update. After the Wrapper is loaded, it checks whether the corresponding resource package has been updated. The Wrapper communicates with the Package Engine to check whether the previously retrieved Package Manifest has been updated. If not updated, then there is no need to refresh the local resources and the Wrapper starts to load the target Web app. Otherwise, if the package has been updated, the Wrapper refreshes the packaged resources stored in the App-Specific Space. Only after the App-Specific Space has finished refreshing, can the target mobile Web app start loading in order to ensure all the resources are up-to-date.

Refreshing Package. When the packaged resources have to be updated, the Wrapper refreshes all the resources together according to the newly retrieved Package Manifest. The update process follows the regular Web resource loading mechanism that each resource request is handled by the Resource Loader of the browsers. For normalized resources, the corresponding concrete resource is requested according to the Resource Mapping.

Loading Resources. When the package updates finishes, the Wrapper loads the target Web app. While loading, the Wrapper intercepts all the resource requests emitted by the Web app. For each resource request, the Wrapper checks the App-Specific Space by matching the URL with regular expressions specified in the Package Manifest. If found, then the resource is directly returned to the app from the App-Specific Space. Otherwise, the request is forwarded to the Resource Loader of browsers to retrieve the corresponding resource either from the cache or from the server.

Note that the update check of resource packages and the refresh of resources are actually handled by the Resource Loader of the client-side browsers to manage the resource package with the regular Web mechanism.

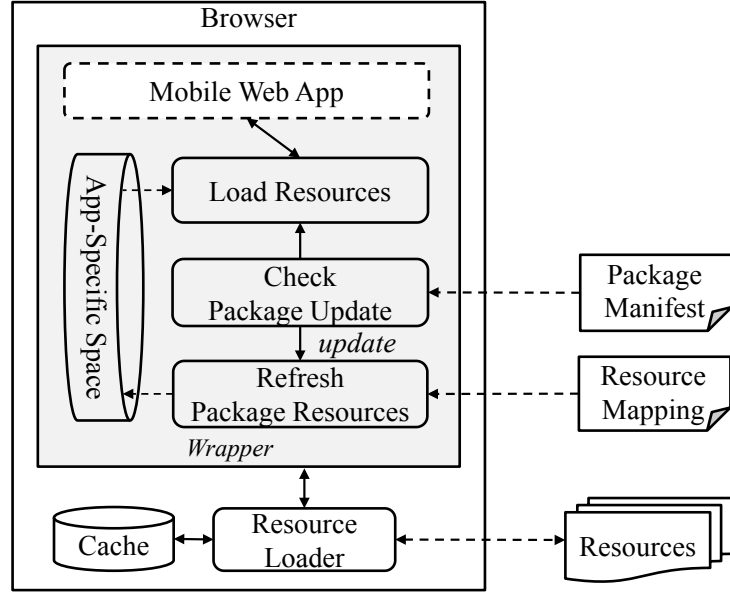


Figure 5: The workflow of the Wrapper.

6 Implementation

We implement ReWAP by leveraging HTML5 [9] Application Cache interface (in short as AppCache). AppCache is designed for offline Web apps, but it actually provides a way for developers to control their resources. Since HTML5 has already been supported by most of the popular commodity mobile browsers such as Chrome, Safari, and Firefox, mobile Web apps with ReWAP can run directly on the latest mobile browsers. Therefore, we can realize the easy and fast deployment without introducing extra cost to end users.

In this section, we first give some background knowledge of HTML5 Application Cache and then describe the details of our implementation and deployment.

6.1 HTML5 Application Cache

The Application Cache (AppCache) is a feature of HTML5 that aims to allow Web apps to be reliably accessed when the browser is offline. To enable AppCache, developers provide a manifest file to specify what resources are needed for Web apps to work offline, and configure the manifest file to the “*manifest*” attribute of an HTML document’s `<html>` tag. The manifest file mainly consists of two sections. Each section has a list of URLs specifying the behavior of the corresponding resources.

- **CACHE.** Resources listed in this section are explicitly cached after they are downloaded for the first time. Even when the browser is online, these resources are still loaded from the AppCache rather than being downloaded from the network. Note that HTML documents referring to manifest files are set in the CACHE section by default.
- **NETWORK.** Resources listed under this header can bypass the AppCache and be requested regularly by the browser. When the browser is offline, these resources cannot be loaded from the AppCache. A wildcard flag “*” can be used to make any resource that is not listed in the CACHE section bypass the AppCache mechanism.

6.2 Implementation of the Package Engine

Based on the AppCache specification, the concept of Package Manifest in ReWAP can be implemented by the manifest file of AppCache. However, as the AppCache manifest can use only concrete URLs, we put the concrete resources in the manifest and use the generated Resource Mapping file to help check whether a resource is in the package or not by matching with the Resource Mapping. We implement the Package Engine in Java. The Package Engine can be published as a stand-alone component and deployed as a service at the server-side. For the browser facility of the Package Engine, we use the Chromium Embedded Framework [10] to render the mobile Web apps and record the HTTP traffic.

6.3 Implementation of the Wrapper

The Wrapper is totally implemented as an HTML page with AppCache enabled together with a JavaScript library. The Wrapper is totally implemented by standard Web technologies so that it can run directly on modern mobile browsers.

The HTML page is configured to use the manifest file generated by the Package Engine. The App-Specific Space and Check Package Update can be provided by the AppCache. To realize loading the target Web app, the page registered a JavaScript callback function on the `onload` event. When loading the HTML page is finished, the callback function is executed to dynamically fetch the root HTML of the target Web app and modify the DOM tree to render the actual page. Therefore, we can ensure that the retrieved HTML file of the target Web app is up-to-date.

To intercept resource requests, we use JavaScript reflection mechanism to register callback functions for all the cases of resource request. When the requested URL matches a URL's regular expression in the Resource Mapping, we replace the requested URL with the corresponding concrete URL to make the resource loaded from the AppCache.

Since the AppCache can work only in the next load after refresh, we explicitly call the `swap()` function of the AppCache when the AppCache's *update* event is triggered in order to make the AppCache use the latest resources.

6.4 Deployment

Given the implementations based on the AppCache, developers can easily deploy ReWAP on their mobile Web architecture, requiring no extra cost to end users.

The whole ReWAP is deployed as a separate service on the Web server. Developers can configure the target mobile Web app to be integrated with ReWAP to launch a certain instance of ReWAP. When the ReWAP is launched, the Package Engine is automatically started as a background process on the server, while the Wrapper specific to the configured app is created in a certain folder. While the Package Engine is running, the two configuration files, Package Manifest and Resource Mapping, are also generated in the same folder as the Wrapper. Developers should make the folder accessible by the standard HTTP protocol where the Wrapper, Package Manifest and Resource Mapping are assigned dedicated URLs.

To make the Wrapper work for the target mobile Web app, the developer needs to only configure the server to redirect the entrance of the Web app to the URL of the Wrapper. Therefore, when users visit the ReWAP-enabled app, the Wrapper is first loaded to the browser. Nevertheless, the end-users are unaware of the existence of ReWAP when they request the target Web apps.

7 Evaluations

We evaluate ReWAP from three main aspects². First, we investigate the overall performance of ReWAP by measuring how much data traffic can be saved for mobile Web apps with ReWAP compared to the original apps that are with regular cache mechanisms enabled. Second, we evaluate the performance of

²Please visit <http://taoxie.cs.illinois.edu/rewap/> to get more information of ReWAP implementation and evaluation results.




Figure 6: Distribution of saved data traffic comparing the case w/ and w/o ReWAP.

the Package Engine, such as the resource normalization and the prediction of update time. Third, we evaluate the overhead of the Wrapper incurred to the mobile Web apps.

7.1 Overall Performance

The performance of ReWAP is measured by how much data traffic can be saved compared to the original mobile Web apps. There are three main factors influencing the performance.

Redundant transfers of the original mobile Web apps. ReWAP aims to reduce redundant transfers for mobile Web apps. Therefore, if a mobile Web app has few redundant transfers, there may be not so much room for ReWAP to contribute to improve.

User revisiting distribution. Since ReWAP generates a unique resource package for all the users of a mobile Web app, we use a user revisiting distribution to calculate the benefit brought by different resource packages and choose the one with the largest benefit in our algorithm of generating resource packages (Algorithm 4.3). Certainly, different distributions could lead to different performances of ReWAP.

Retrieving interval. ReWAP helps reduce redundant transfers for retrieving resources of the same mobile Web apps. Generally, the longer the revisiting interval is, the lower the performance of ReWAP is because more resources may change after a longer interval, leading to the invalid of resource packages.

Given the preceding factors, we apply a simulation-based approach to evaluate the performance of ReWAP. We chose 50 mobile Web apps that suffer heavily from redundant transfers according to our previous measurement studies in WWW 2015 work [4]. Then we recorded all the resources of each mobile Web app’s homepage every 30 minutes and lasted for 15 days. We assume that 30-minutes are sufficiently short to capture the changes of stable resources of Web apps. Based on this data set, we use the first five days’ record to train our algorithms of normalization (Algorithm 4.1) and prediction (Algorithm 4.2), and use the last ten days’ records to generate resource packages. We evaluate the performance of ReWAP for revisiting intervals ranging from 0.5 hour to one day with 0.5 hour as an interval. Therefore, we have 48 revisiting intervals to investigate (0.5 hour, 1 hour, 1.5 hour, ..., 24 hours). For user revisiting distributions, given a revisiting interval to investigate, we assume a 100% revisiting of the same interval as we are to investigate. For comparison, we simulate an ideal cache with unlimited size to visit each Web app by each given revisiting interval. We calculate the differences of data traffic consumed by mobile Web apps with and without ReWAP. Note that the data traffic of ReWAP includes the cost of the update of resource packages.

Figure 6 shows the overall performance of ReWAP for different revisiting intervals. We draw a box plot to represent the distribution of saved data traffic for each revisiting interval. The median value of




Figure 7: Distribution of the resource coverage among different intervals.

saved data traffic varies from 8% to 51%, indicating that mobile Web apps with ReWAP can reduce averagely 8% to 51% of the data traffic compared to the original Web apps with cache enabled. When the revisiting interval becomes larger, the saved data traffic decreases because the resource package has to be refreshed due to the update of packaged resources.

The variance of saved data traffic is significantly large for all revisiting intervals. In the best cases, the saved data traffic can reach almost 100%, such as revisiting intervals shorter than 4 hours, implying that almost all the resources that should be downloaded from the network by the original Web app can be directly loaded from the local storage by ReWAP. In some worst cases especially for the longer revisiting interval, there is hardly improvement of achieved ReWAP because resources of some mobile Web apps are not stable enough to be packaged for longer revisiting intervals so the resources always change between two visits.

We should mention that more data traffic can be saved by ReWAP in the practice where the size of local cache is limited. The evaluation presented in this paper assumes an ideal cache where cached resources are always kept in the cache. However, in the real case, the size of cache is limited on mobile devices and cached resources are often removed out of the cache. Since ReWAP maintains an app-specific space for each Web app, the removed resources of original Web apps are also likely to be loaded from the local storage.

7.2 Performance of the Package Engine

The performance of ReWAP is determined by the resource packages generated by the Package Engine. More captured resources, more accurate predicted update time, and more stable resource packages can lead to saving more data traffic. We evaluate the performance of Package Engine by the intermediate data gathered during the experiment in Section 7.1.

The Package Engine maintains a list of historic resources that are the candidates to be packaged. The list is changing every time when the Resource Retriever retrieves the new status of resources. It can be better if the resource list at a certain time t covers more resources at the time later than t . Figure 7 shows the distribution of the rate of resource coverage after different intervals ranging from 0.5 hour to one day. We can observe that the median coverage rate is around 70% and it is very stable for different revisiting intervals. This result mainly accounts for the contribution of our normalization technique that cross-dress resources can be normalized to one resource. Therefore, more resources can be covered for longer durations. The Package Engine uses an estimated update time to judge whether a resource is stable. Since the update time is an important factor in calculating the benefit of resource




Figure 8: Distribution of the accuracy of predicted update time among different intervals.

packages in Algorithm 4.2, the predictions need to be precise enough. Figure 8 shows the distribution of the precision of predicted update time. We can observe that the precision decreases as the interval increases. For all the intervals, the median predicting precision is above 85%. For intervals less than 5 hours, the median precision is 100%. Overall, such accuracy can be satisfying to most apps and demonstrates the effectiveness of ReWAP. The more stable the resource packages are, the less refresh the Wrapper performs. We calculate the time duration between every two updates of resource packages. Figure 9 shows the distribution of the duration. We can observe that the median duration of a resource package is 5 hours, indicating that resource packages should be updated every 5 hours in the medium cases. Therefore, the performance of ReWAP is better for revisiting intervals less than 5 hours in the situation of the experiment in Section 7.1.

7.3 Overhead of the Wrapper

When a ReWAP-enabled mobile Web app runs on the browsers, the Wrapper can introduce overhead compared to the original app. The overhead lies in two main aspects. One is the manifest files to specify resource packages. The other is the computation logics of resource mapping and the Application Cache itself.

Figure 10 shows the distribution of manifest files’ size, which is gathered from the manifests generated in the experiment of Section 7.1. We can see that the median size is only 5 KB and the largest is not more than 20 KB. Therefore, the overhead of manifest file is small enough.

To evaluate the overhead of computation logics and their influences on the page load time, we generate some test pages whose number of resources ranges from 20 to 100, and each resource is 100 KB. We assume that all the resources of each page are put into the corresponding package and 10% of the packaged resources are normalized resources that are identified by URL regular expressions. With the assumption, we generate the manifest files. Then we visit each page twice in the browser on a smartphone of Samsung Galaxy S4 with Android 5.0. We record the CPU usage and memory usage as well as the page load time during loading each page with and without ReWAP. We find that the average CPU usage is increased by 15% for pages with ReWAP while the memory usage is of no significant differences.

Figure 11 shows the page load time for different pages in four cases: cold start without ReWAP, cold start with ReWAP, warm load without ReWAP, and warm load with ReWAP. We can observe that as the number of resources increases, the page load time increases in all cases, and the gap between cold start and warm load becomes bigger no matter whether the app is equipped with or without ReWAP. For the cold start, the page load time of pages with ReWAP is a little longer than those without ReWAP.




Figure 9: Distribution of the durations of resource package.




Figure 10: Distribution of the size of Package Manifest.




Figure 11: Influences of ReWAP on Page load time for cold start and warm load.

However, for the warm load, the page time of pages with ReWAP is much shorter than those without ReWAP. This observation implies that ReWAP can also reduce the page load time when mobile Web apps are revisited. In the best cases, the page load time is reduced by more than 50%.

In summary, we can conclude that the overhead of the Wrapper is very minor. ReWAP can even reduce the page load time when mobile Web apps are revisited. Therefore, the results demonstrate that ReWAP is practical to be adopted and applied.

8 Related Work

It is well known that the user experiences of mobile Web apps are far from satisfaction in terms of the page load time, data traffic drain, and energy consumption. Redundant transfers of resources are the most dominant issue leading to such inefficiency. Our work focuses on reducing redundant transfers to improve the user experiences of mobile Web apps. We then discuss related work.

Measurement studies on resource loadings of mobile Web apps. Wang et al. [11] advocated that resource loading contributes most to the browser delay. Wang et al. [3] designed a lightweight in-browser profiler, called WProf, and studied the dependencies of activities when browsers load a webpage. Nejati et al. [12] extended WProf to WProf-M and studied the differences of page loading process between mobile and non-mobile browsers. Li et al. [13] designed WebProphet to capture dependencies among Web resources and to automate the prediction of user-perceived Web performance. The poor performance of mobile Web cache is a key issue leading to redundant transfers. Qian et al. [2] measured the performance of mobile Web cache in terms of the cache implementation and revealed that about 20% of the total Web traffic examined is redundant due to imperfect cache implementations. In their later work [14], they studied the caching efficiency for the most popular 500 websites and found that caching is poorly utilized for many mobile sites. Wang et al. [7] found that cache has very limited effectiveness: 60% of the requested resources are either expired or not in the cache. Our previous work [4] adopted a proactive approach to measure the performance of mobile Web cache and found that more than 50% of resource requests are redundant on average for 55 popular mobile websites. In particular, we found some underlying factors leading to redundant transfers of mobile Web apps, i.e., Same Content, Heuristic Expiration, and Conservative Expiration Time. These measurement studies motivate us to reduce redundant transfers for mobile Web apps.

Techniques to reduce redundant transfers. Wang et al. [15] investigated how Web browsing can benefit from micro-cache that separately caches layout, code, and data at a fine granularity. They studied

how and when these resources are updated, and found that the layout and code that block subsequent object loads are highly cacheable. Our resource packaging can be viewed as to realize similar features proposed by micro-cache in the resource granularity. Most of the stable layout and code resources are put into the package to always be loaded from the local environment. Zhang et al. [16] implemented a system-wide service called CacheKeeper, to effectively reduce overhead caused by poor Web caching of mobile apps. CacheKeeper can also work for browsers but it relies on the support of operating systems. Our implementation of ReWAP utilizes the standard HTML5, which has been supported by all the commodity mobile Web browsers. So we can achieve fast and easy deployment with no cost to end users.

General techniques to improve the performance of mobile Web. Some previous work focuses on improving the compute-intensive operations for mobile Web browsers, such as style formatting [17], layout calculation [18, 19], and JavaScript execution [20, 21]. However, Wang et al. [11] argued that the key to improve the performance of mobile Web is to speed up resource loading. Various solutions have been proposed to optimize resource loading. These solutions include new network protocols such as SPDY [22] and HTTP2 [23], browser optimization such as prefetching [24] and speculative loading [7], and proxy-based systems such as Flywheel [25] and KLOTSKI [26]. These previous solutions are orthogonal to reducing redundant transfers.

Dynamics and user revisits of Web pages. Douglass et al. [27] performed a live study on the influences of resource changes and user revisits on the Web caching in early 1997. Fetterly et al. [28] measured the degree of Web page changes and investigated the factors correlated with change intensity. Adar et al. [29] studied the Web revisiting behaviors from a live data set. They identified four revisiting patterns for different kinds of Web pages. In their later work [30], they studied the relationship between the dynamics of Web pages and user revisiting patterns. Although all these previous efforts focus on the Web for desktop computers, their findings can be partly leveraged by the mobile Web. Our work depends on the dynamics and user revisits of Web pages to maintain the resource package.

9 Conclusion

In this paper, we have presented the ReWAP approach, by radically making the resource management of mobile Web apps perform in a similar fashion to the native apps. The key rationale of ReWAP is to provide more efficient and “application-aware” control of resource management rather than relying on only the current mechanisms such as Web cache, to avoid the caused unnecessary redundant resource transfer. To realize the efficient resource packaging, our ReWAP includes a normalization technique to identify the same resources but with different URLs, a learning-based technique to accurately predict the updates of resources, and an algorithm to minimize the refresh frequency of resource packages as well as reduce the overhead. We have evaluated ReWAP based on long-term (15-day) traces of existing Web apps that suffer from redundant resource transfer, and the results demonstrate our approach’s effectiveness and efficiency.

Given that ReWAP can be easily deployed into existing Web applications with very few manual efforts, our ongoing work is to encapsulate ReWAP as an independent module into currently popular Web servers such as Apache, Nginx, and Node.js. When ReWAP is deployed and the actual access logs are obtained, we can derive the distribution of user’s visit frequency for a given Web app, and thus optimize the prediction algorithm by tuning the parameters and designing online learning kernels.

While our work in this paper addresses the key challenges caused by the dynamics of Web apps, some other issues are worth further exploring. For example, although the data traffic can be significantly reduced, it is observed that the page load time of Web apps does not always improve due to the checking update with server-side components. Prefetching the updates of resource packages at a desirable time or the push-oriented update notification can be potential solutions. Other performance issues, e.g., the energy drain, need to be addressed as well.

References

- [1] N. Serrano, J. Hernantes, and G. Gallardo, “Mobile web apps,” *Software, IEEE*, vol. 30, no. 5, pp. 22–27, 2013.
- [2] F. Qian, K. S. Quah, J. Huang, J. Eрман, A. Gerber, Z. Mao, S. Sen, and O. Spatscheck, “Web caching on smartphones: ideal vs. reality,” in *Proceedings of the 10th international conference on Mobile systems, applications, and services, MobiSys 2012*, 2012, pp. 127–140.
- [3] X. S. Wang, A. Balasubramanian, A. Krishnamurthy, and D. Wetherall, “Demystifying page load performance with WProf,” in *Proceedings of USENIX Conference on Networked Systems Design and Implementation, NSDI 2013*, 2013, pp. 473–485.
- [4] Y. Ma, X. Liu, S. Zhang, R. Xiang, Y. Liu, and T. Xie, “Measurement and analysis of mobile web cache performance,” in *Proceedings of the 24th International Conference on World Wide Web, WWW 2016*, 2015, pp. 691–701.
- [5] X. Liu, Y. Ma, Y. Liu, T. Xie, and G. Huang, “Demystifying the imperfect client-side cache performance of mobile web browsing,” *IEEE Transactions on Mobile Computing*, 2015, online Accessed.
- [6] “Rfc 2616.” [Online]. Available: <http://www.w3.org/Protocols/rfc2616/rfc2616.txt>
- [7] Z. Wang, F. X. Lin, L. Zhong, and M. Chishtie, “How far can client-only solutions go for mobile browser speed?” in *Proceedings of the 21st international conference on World Wide Web, WWW 2012*, 2012, pp. 31–40.
- [8] T. Zhang, “Solving large scale linear prediction problems using stochastic gradient descent algorithms,” in *Proceedings of the twenty-first international conference on Machine learning (ICML)*, 2004, pp. 919–926.
- [9] “HTML5.” [Online]. Available: <http://www.w3.org/TR/html5/>
- [10] “Chromium embedded framework.” [Online]. Available: <https://bitbucket.org/chromiumembedded/cef/>
- [11] Z. Wang, F. X. Lin, L. Zhong, and M. Chishtie, “Why are web browsers slow on smartphones?” in *Proceedings of the 12th Workshop on Mobile Computing Systems and Applications, HotMobile 2011*, 2011, pp. 91–96.
- [12] J. Nejati and A. Balasubramanian, “An in-depth study of mobile browser performance,” in *Proceedings of the 25th International Conference on World Wide Web, WWW 2016*, 2016.
- [13] Z. Li, M. Zhang, Z. Zhu, Y. Chen, A. Greenberg, and Y.-M. Wang, “Webprophet: Automating performance prediction for web services,” in *Proceedings of the 7th USENIX conference on Networked systems design and implementation, NSDI 2010*, 2010, pp. 143–158.
- [14] F. Qian, S. Sen, and O. Spatscheck, “Characterizing resource usage for mobile web browsing,” in *Proceedings of the 12th annual international conference on Mobile systems, applications, and services, MobiSys 2014*, 2014, pp. 218–231.
- [15] X. S. Wang, A. Krishnamurthy, and D. Wetherall, “How much can we micro-cache web pages?” in *Proceedings of the 2014 Internet Measurement Conference, IMC 2014*, 2014, pp. 249–256.
- [16] Y. Zhang, C. Tan, and L. Qun, “Cachekeeper: A system-wide web caching service for smartphones,” in *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp 2013*, 2013, pp. 265–274.

- [17] H. Wang, M. Liu, Y. Guo, and X. Chen, “Similarity-based web browser optimization,” in *Proceedings of the 23rd International World Wide Web Conference, WWW 2014*, 2014, pp. 575–584.
- [18] K. Zhang, L. Wang, A. Pan, and B. B. Zhu, “Smart caching for web browsers,” in *Proceedings of the 19th International Conference on World Wide Web, WWW 2010*, 2010, pp. 491–500.
- [19] D. Mazinanian, N. Tsantalis, and A. Mesbah, “Discovering refactoring opportunities in cascading style sheets,” in *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, 2014, pp. 496–506.
- [20] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl, “Anatomizing application performance differences on smartphones,” in *Proceedings of the 8th International Conference on Mobile Systems, Applications, and Services, MobiSys 2010*, 2010, pp. 165–178.
- [21] L. Gong, M. Pradel, and K. Sen, “JITProf: pinpointing jit-unfriendly javascript code,” in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015*, 2015, pp. 357–368.
- [22] X. S. Wang, A. Balasubramanian, A. Krishnamurthy, and D. Wetherall, “How speedy is spdy?” in *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2014*, 2014, pp. 387–399.
- [23] “Http/2.” [Online]. Available: <https://http2.github.io/>
- [24] D. Lymberopoulos, O. Riva, K. Strauss, A. Mittal, and A. Ntoulas, “Pocketweb: instant web browsing for mobile devices,” in *Proceedings of the 17th International Conference on Architectural Support for Programming Languages and Operating Systems, APSLOS 2012*, 2012, pp. 1–12.
- [25] V. Agababov, M. Buettner, V. Chudnovsky, M. Cogan, B. Greenstein, S. McDaniel, M. Piatek, C. Scott, M. Welsh, and B. Yin, “Flywheel: Google’s data compression proxy for the mobile web,” in *Proceedings of 12th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2015*, 2015, pp. 367–380.
- [26] M. Butkiewicz, D. Wang, Z. Wu, H. V. Madhyastha, and V. Sekar, “Klotski: Reprioritizing web content to improve user experience on mobile devices,” in *12th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2015*, 2015, pp. 439–453.
- [27] F. Douglass, A. Feldmann, B. Krishnamurthy, and J. C. Mogul, “Rate of change and other metrics: a live study of the world wide web,” in *Proceedings of the 1st USENIX Symposium on Internet Technologies and Systems, USITS’97*, 1997.
- [28] D. Fetterly, M. Manasse, M. Najork, and J. L. Wiener, “A large-scale study of the evolution of web pages,” in *Proceedings of the Twelfth International World Wide Web Conference, WWW 2003*, 2003, pp. 669–678.
- [29] E. Adar, J. Teevan, and S. T. Dumais, “Large scale analysis of web revisitation patterns,” in *Proceedings of the 2008 Conference on Human Factors in Computing Systems, CHI 2008*, 2008, pp. 1197–1206.
- [30] —, “Resonance on the web: web dynamics and revisitation patterns,” in *Proceedings of the 27th International Conference on Human Factors in Computing Systems, CHI 2009*, 2009, pp. 1381–1390.